



Coherent Naming Schemes: A Case Study


Matthew F. Ringel
Tufts University
ringel@net.tufts.edu






Overview

- In early 2004, the Tufts NOC recognized the need for a standardized network equipment naming scheme to support:
 - Routers as both entities, and interfaces.
 - Switches
 - Wireless APs
 - Security devices (firewalls, gateways, etc.)
 - Future needs.




The Problem

- 3 campuses, 15 routers, 400+ switches, and many different existing name formats.
- Previously there was no standard for device naming.
- Names were chosen in an attempt to convey useful information, but a lack of standards had deleterious effects on productivity, complicating troubleshooting, equipment deployment, documentation, and automation efforts.




What didn't work: (part 1)

- Inconsistent nomenclature for buildings
 - Postal addresses mixed with building common names and University building ID numbers.
 - Location information by floor or nearby departments rather than room number.
- Inconsistent data in name
 - TAB-DC-15K-A1
 - TAB-BI
 - Research-2-FI8K
- Arbitrary interface names
 - sack-rtr-5
 - tufts-pri-border




What didn't work: (part 2)

- Relying on institutional memory
 - “Everyone knows that the switch for Cousens Gym is actually in Halligan Hall in that room that doesn't have a number on the door.”
- Names based on device hardware
 - Sackler-7507
- The myth of “naming security”
 - Obscuring the names of network devices in the name of security isn't as important as being able to do effective troubleshooting and documentation.
 - Sufficiently dedicated attacker etc. etc.




What we needed:

- Comprehensible – Human readable with immediate meaning.
- Extensible – Must accommodate future device types.
- Derivable – Formulaic. Completely predictable names given a set of parameters.
- Self-Documenting – Names defined by role. i.e. Traceroute output should make sense.
- Unique – One name => One Device.




Some Counterexamples

- Using TelCo codes to name locations
 - cmbrmaks, somrmats (??)
 - May not have the resolution or information you need, and **will** become inconsistent with whoever you patterned yourself on.
- Compression artifacts
 - Everyone abbreviates differently, names become unpredictable and underivable.
- Avoid using equipment types in names
 - Equipment is replaced or upgraded, should the name also need to be replaced?
 - The name should reflect a location and role, not a specific piece of hardware. This keeps name churn to a minimum.




Methodology

- Use authoritative naming sources
 - Space planning dept. for building names.
 - Building and Grounds Dept. for actual room numbers.
- Formal grammar
 - Enforce a strict structure to names.
 - Human parsing now, but computer parsing in the near future.
- List assumptions about the network
 - Combined L2/L3 core with MAN/WAN links to other campuses
 - Three basic kinds of links
 - What are you optimizing for?




Methodology (part 2)

- Informative is more important than brief
 - Want brief? Use a CNAME.
 - People will type in long names if the names:
 - Make sense.
 - Can be derived correctly on the first try.
- DNS is not for asset management
 - Never use vendor type, make, or model.




The Specifics

- Two syntaxes: multi-homed devices (routers) and single-homed devices (switches, APs, etc.)
- Routers
 - Three types of links
 - b = border (i.e. administrative handoff)
 - x = transit (i.e. interim hop, usually no clients)
 - t = terminus (i.e. last hop, usually a client network or loopback)
 - Router names: (role)-priority
 - E.g.: sackler-rtr-pri, grafton-rtr-pri, border-sec
 - Why pri and sec? Why not 01 and 02?
 - Devices are actually primary and secondary.




The Specifics (part 2)

- Name L3 links from origin point, forward
 - E.g. the L3 interface of the Sackler router connecting to the Grafton router is:
 - sackler-rtr-pri-x-grafton-rtr-pri
 - Likewise the L3 interface of the Grafton router at the other end of the same link is:
 - grafton-rtr-pri-x-sackler-rtr-pri
 - i/f Vlan80 (a leaf subnet) on the Anderson hub router:
 - anderson-rtr-pri-t-vlan80
 - Handoff to the standalone EECS dept. network, from the Anderson router:
 - anderson-rtr-pri-b-eeecs




The Specifics (part 3)

- Switches and other single-homed devices (i.e. WAPs)
 - Single “interface”, simpler name
 - <building name>-rm<room number>-<type><instance>
 - The first switch in room 124 of the 4 Colby St. building is:
 - 4colby-st-rm124-unit0
 - The first WAP in the ceiling of room 168 in Paige Hall is:
 - paige-hl-rmc168-ap0




Implementation

- Several days of planning, and pre-generating the names beforehand
- No concessions were made at runtime
 - If it's hard to use, we had to know immediately.
- Daytime changes
 - In-group coordination only, in order to change over monitoring systems where necessary.
- Backward compatibility
 - CNAMEs were put in place for names already in use by “finger macros”.




The Results: Success!

- A consistent, derivable, and self-documenting naming scheme.
- Faster network debugging times.
- Faster security response times.
- Faster turnaround on addition of new devices.
- People **did** adjust quickly and didn't mind typing in longer names.
 - Typing a longer name is still faster than needing to look it up in an outside reference.



Lessons Learned

- You'll always have to make assumptions about network topology
 - E.g. more meaningful to name a point-to-point as its endpoints than as "t" interface
 - "universal" is hard, and rarely as useful as it sounds (q.v. X.500)
- Name everything first
 - Pre-deriving all current names was the best test we could really do of the naming scheme's flexibility



Acknowledgements

- From the NOC Group at tufts.edu
 - Marc Jimenez, Robin Garner, Bob Moran
 - Joel Gridley, Linda Van Horn
 - Peter Radcliffe
- Others
 - Nicolai Plum, Cat Okita, Aaron Block